

ABSTRACT

The present invention provides a method of estimating cardinalities using statistical soft constraints in a relational database management system. The statistical soft constraints are generated from statistical information gathered with respect to virtual columns. The virtual columns are the result of operations involving one or more columns in the database. When a query is received, the database management system determines whether any of the virtual columns are of interest because they involve expressions or relations that are also reflected in the query. The database management system then develops statistical soft constraints from the statistical characteristics of the virtual column and uses those statistical soft constraints to assist in estimating the cardinality of a step in the query plan. The virtual columns may or may not be actually materialized into a column of values.